XaasIO Software-Defined Storage

Block and Object Storage, Built for Private Cloud, Kubernetes and AI

Give private cloud, virtualization, Kubernetes, backup and AI workloads a shared storage foundation. Deliver persistent block volumes through Ceph RBD and S3-compatible object storage through Ceph Object Gateway, on infrastructure you control.



© 2026 XaasIO. All rights reserved. | v2.0

ONE PLATFORM FOR BLOCK AND OBJECT STORAGE

Consume storage through familiar cloud, virtualization and application interfaces.

VOLUMES AND BUCKETS

Expose RBD-backed volumes and S3-compatible buckets through supported application and platform integrations.

PRIVATE AND AIR-GAPPED

Keep storage services, identity, monitoring and approved software updates within your security boundary.

INDEPENDENT SDS

Operate storage as a dedicated platform. Add XaasIO integrations and operational services as workloads evolve.

CORE STORAGE CAPABILITIES

Persistent Block Storage

Create, attach, resize, snapshot and clone Ceph RBD volumes. Connect OpenStack Cinder, Nova and Glance, QEMU/libvirt and Kubernetes CSI through supported integrations.

S3-Compatible Object Storage

Manage buckets, policies, quotas, versioning, lifecycle rules and supported Object Lock configurations. Validate the S3 API operations required by each application.

Placement, Protection and Media

Use replication or erasure coding with CRUSH placement across chosen failure domains. Match HDD, SSD and NVMe device classes to capacity, latency and endurance needs.

Identity, Encryption and Governance

Use CephX and RGW access controls. Design TLS, device encryption, RBD image encryption, object encryption and external key management for the selected deployment.

Recovery, Backup and DR

Use checksums, scrubbing and recovery for local resilience. Scope asynchronous RBD mirroring and RGW multisite separately from independent backup and application recovery.

Management and Observability

Monitor health, capacity, device condition, latency and recovery with Ceph Manager and Dashboard. Add XaasIO MLT for centralized monitoring and operational visibility.

DEPLOYMENT ARCHITECTURE

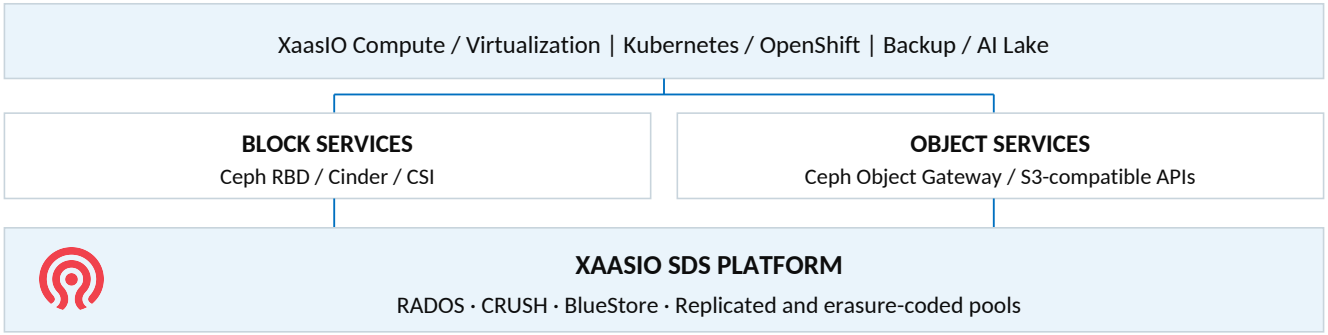

DEPLOYMENT CHOICES

Dedicated storage, converged cloud storage or hyperconverged infrastructure. Use Rook-managed Ceph where required by the Kubernetes deployment design.

OPTIONAL EXTENSIONS

Add XaasIO CMP, MLT or Hyperscaler integrations. Scope NVMe-oF, CephFS/NFS, migration and DR services separately. Core storage can operate independently of these extensions.

© 2026 XaasIO. All rights reserved. | v2.0

42U SAMPLE RACK SOLUTION

Dedicated Ceph storage for capacity-oriented block and object workloads.

16	192	3.84 PB
Storage nodes	20 TB data drives	Raw data capacity

42U FRONT ELEVATION

42-41	Reserved / airflow
40	OOB management
39	100 GbE fabric B
38	100 GbE fabric A
37	RGW gateway 02
36	RGW gateway 01
35	Control 03
34	Control 02
33	Control 01
32-31	SDS-16 | 12 × 20 TB
30-29	SDS-15 | 12 × 20 TB
28-27	SDS-14 | 12 × 20 TB
26-25	SDS-13 | 12 × 20 TB
24-23	SDS-12 | 12 × 20 TB
22-21	SDS-11 | 12 × 20 TB
20-19	SDS-10 | 12 × 20 TB
18-17	SDS-09 | 12 × 20 TB
16-15	SDS-08 | 12 × 20 TB
14-13	SDS-07 | 12 × 20 TB
12-11	SDS-06 | 12 × 20 TB
10-09	SDS-05 | 12 × 20 TB
08-07	SDS-04 | 12 × 20 TB
06-05	SDS-03 | 12 × 20 TB
04-03	SDS-02 | 12 × 20 TB
02-01	SDS-01 | 12 × 20 TB

A/B 0U PDUs · Dual server power supplies

ILLUSTRATIVE HARDWARE

EQUIPMENT	QTY / U	EXAMPLE SPECIFICATION
Storage	16 × 2U	Per node: 12 × 20 TB enterprise CMR HDD; 32 physical cores; 256 GB ECC RAM; 2 × 100 GbE.
OS + DB/WAL	In nodes	Per node: mirrored 2 × 960 GB boot SSD; 2 × 7.68 TB enterprise NVMe for BlueStore DB/WAL. Excluded from data capacity.
Control	3 × 1U	16 cores; 64 GB RAM; mirrored SSD. Three MONs and active/standby MGR services.
Object gateways	2 × 1U	16 cores; 64 GB RAM; mirrored SSD; 2 × 100 GbE. RGW and redundant ingress.
Data network	2 × 1U	Two 32-port 100 GbE switches. Dual-attached servers; segregated client and recovery traffic.
Management	1 × 1U	48-port out-of-band switch.
Reserved	2U	Cabling and airflow allowance. Total: 42U.

TRANSPARENT CAPACITY MATH

Raw: 16 × 12 × 20 TB = **3,840 TB**
Protected: raw ÷ 3, × 4/6, or × 8/11
Planning budget: protected capacity × 70%

Use direct-access/JBOD for data OSD drives. Validate chassis slots, endurance and DB/WAL sharing. Size A/B power feeds, cooling and floor loading to the final BOM; UPS is external. This HDD capacity example is not an all-flash performance tier.

WHOLE-RACK PROTECTION ALTERNATIVES

PROTECTION	LAYOUT	EFFICIENCY	PROTECTED CAPACITY	AT 70% FILL
3× replication	3 full copies	33.3%	1,280 TB	896 TB
EC 4+2	4 data + 2 parity	66.7%	2,560 TB	1,792 TB
EC 8+3	8 data + 3 parity	72.7%	2,792.7 TB	1,954.9 TB

Alternative calculations over the same 3,840 TB; do not add them. A 70% fill target reserves 30% for operating and recovery headroom. Decimal units: 1 PB = 1,000 TB. Figures exclude metadata, indexes, formatting and EC padding; no compression or deduplication gain is assumed.

© 2026 XaasIO. All rights reserved. | v2.0

POOL DESIGN, LIFECYCLE AND SUPPORT

Match storage protection to the workload and govern changes across the lifecycle.

ONE RACK WITH THREE COEXISTING POOL GROUPS

Example allocation of the same 3,840 TB raw capacity, using the 70% fill target.

POOL / INTENDED USE	RAW ALLOCATION	LOGICAL BUDGET*
3× replicated: block + replicated metadata	20% / 768 TB	179.2 TB
EC 4+2: object, backup and data-lake payloads	40% / 1,536 TB	716.8 TB
EC 8+3: large-object capacity payloads	40% / 1,536 TB	782.0 TB

*Combined budget: approximately 1,678 TB (1.678 PB), before metadata, indexes and allocation overhead. Replicated metadata consumes part of the first pool budget. These are shared-cluster budgets, not disk partitions; enforce quotas and monitor aggregate use.

HOST-DOMAIN RESILIENCE

For 3× / EC 4+2 / EC 8+3, use at least 3 / 6 / 11 distinct storage hosts and proposed min_size 2 / 5 / 9. From a healthy state, per-object recoverability is 2 / 2 / 3 host losses; data-pool I/O continues through 1 / 1 / 2 losses before recovery. Metadata and quorum can limit service sooner.

DEPLOYMENT GUARDRAILS

A single rack does not provide rack or site resilience. Keep RGW index/metadata and RBD metadata replicated. RBD data on EC requires BlueStore and allow_ec_overwrites. Benchmark EC 8+3 writes and recovery; validate failure scenarios, headroom and independent backup/DR before production.

PREDICTABLE RELEASES AND CONTROLLED LIFECYCLE

26.06	25.12	26.12	27.12
Nammalvar	Valluvar	Ilango Adigal	Avvaiyar
Current standard	Current LTS baseline	Planned standard	Planned LTS

Portfolio milestones: June and December releases, YY.MM versions and YY.MM.N maintenance updates. LTS designations occur every 24 months; this does not define the support term. Confirm SDS availability, support duration, EOL and upgrade paths in the release record and support agreement.

RELEASE EVIDENCE

Upstream and XaasIO release SBOMs, artifact verification and provenance are listed per release. Validated air-gapped updates use approved offline bundles, local repositories and retained change records.

SUPPORT AND SERVICES

Architecture, sizing, deployment, acceptance, runbooks, training and lifecycle support. Customer-operated, co-managed, fully managed and dedicated Storage SRE Pod services are available.

XAASIO SOFTWARE-DEFINED STORAGE

Build storage around your workloads.

Explore xaasio.com/platform/block-and-object-storage/

Book an architecture workshop at xaasio.com/contact/

AMERICAS
+1 650 523 6628

EMEA
+44 20 3031 1074

APAC
+91 93423 61010

© 2026 XaasIO Group. All rights reserved. Information is subject to change. | xaasio.com